

Space Filling Quantizers for Digital Watermarking

Related Application Data

This patent application claims the benefit of US Provisional application 60/256,629, filed December 18, 2000, which is hereby incorporated by reference.

5

Technical Field

The invention relates to steganography, data hiding, and digital watermarking.

Background and Summary

Digital watermarking is a process for modifying physical or electronic media to embed a machine-readable code into the media. The media may be modified such that the embedded code is imperceptible or nearly imperceptible to the user, yet may be detected through an automated detection process. Most commonly, digital watermarking is applied to media signals such as images, audio signals, and video signals. However, it may also be applied to other types of media objects, including documents (e.g., through line, word or character shifting), software, multi-dimensional graphics models, and surface textures of objects.

Digital watermarking systems typically have two primary components: an encoder that embeds the watermark in a host media signal, and a decoder that detects and reads the embedded watermark from a signal suspected of containing a watermark (a suspect signal). The encoder embeds a watermark by altering the host media signal. The reading component analyzes a suspect signal to detect whether a watermark is present. In applications where the watermark encodes information, the reader extracts this information from the detected watermark.

Several particular watermarking techniques have been developed. The reader is presumed to be familiar with the literature in this field. Particular techniques for embedding and detecting imperceptible watermarks in media signals are detailed in the assignee's co-pending application serial number 09/503,881 and US Patent 5,862,260,

which are hereby incorporated by reference. A form of watermark embedding referred to as Quantization Index Modulation (QIM) is described in WO9960514A1, SYSTEM, METHOD, AND PRODUCT FOR INFORMATION EMBEDDING USING AN ENSEMBLE OF NON-INTERSECTING EMBEDDING GENERATORS and

- 5 corresponding US Patent application 09/206,806, filed December 7, 1998; which are hereby incorporated by reference.

This invention provides a method of QIM digital watermark embedding using space filling quantizers. A QIM watermark embedder uses a space filling quantizer to imperceptibly embed an auxiliary message signal into a media signal, such as an image, audio or video signal. A specific example of the space filling quantizer is a hexagonal lattice quantizer.

10 A QIM watermark decoder uses a compatible space filling quantizer to decode the auxiliary message signal from a potentially distorted version of the watermarked media signal.

15

Brief Description of Drawings

Fig. 1 illustrates a block diagram of QIM watermarking.

Fig. 2 illustrates quantizers used for dithered uniform scalar QIM watermarking.

Fig. 3 illustrates an extension of scalar dithered QIM to two dimensions.

- 15 Fig. 4 illustrates an example of a sphere packing arrangement to show how to improve the digital watermark message rate for a given watermark distortion constraint.

Fig. 5 compares the performance of space filling QIM performance with and without distortion compensation.

Fig. 6 shows performance of QIM watermarking with three different quantizer types.

25

Fig. 7 shows performance of distortion compensated QIM watermarking with three different quantizer types.

Detailed Description

INTRODUCTION

5 Digital watermarking is both an active area of research and a growing source of commercial activity. Early research focused upon watermarking as a solution to verify ownership of images and other media, but watermarking has since found its way into a variety of applications. Some uses for digital watermarking include counterfeit prevention, broadcast monitoring, brand protection, and enabling of content on the
10 internet.

In these and other applications, digital watermarking allows both the unobtrusive embedding of information into media as well as the capability to reliably recover the information, possibly after application specific degradation of the watermarked media signal (e.g., audio, image, or video). There is a natural tension between requirements of
15 being able to embed a large amount of information, requirements for the embedded information to imperceptibly alter the original source, and requirements of being able to recover the information after severe distortions to the watermarked source signal. Information theory has been applied to the watermarking problem and has been used to determine the space of requirements over which we can hope to build real systems.

20 Chen and Wornell proposed Quantization index modulation (QIM) as an efficient method of digital watermarking. QIM is used in the information theoretic proofs of watermarking capacity. This capacity can be demonstrated via random coding using a sequence of vector quantizers of increasing block length, with both channel capacity and optimal rate-distortion performance being reached in the limit of infinite quantizer block
25 length. It is tempting to investigate the utility of QIM in practical watermarking situations as well. It is possible that different considerations should be used in the design of quantizers for QIM than for the design of quantizers for compression. Because the quantized values are not digitally encoded in QIM, the number of reconstruction values in each quantizer is not a design constraint, as it is in the design of a conventional quantizer.
30 The lack of a rate constraint in QIM suggests that quantizer design for QIM involves different considerations than does quantizer design for rate-distortion performance.

Lookabaugh has identified three types of advantages of vector quantizers vs. scalar quantizers. See T.D. Lookabaugh and R.M. Gray, "High-Resolution Quantization Theory and the Vector Quantizer Advantage," IEEE Trans. On Information Theory 35, pp. 1020-1033, Sept. 1989. These advantages are called the space-filling, shape, and memory advantages. The space-filling advantage is due to the increasing efficiency with which polytopes can be used to fill space as the dimension increases. The shape advantage involves the marginal distribution of the source signal and the tendency for the typical source distribution to be concentrated in a small volume in high dimensional space. Finally, the memory advantage exploits dependencies between samples of the source signal. This document describes digital watermarking methods that exploit the space-filling advantage in particular, and provides a heuristic argument that the space-filling advantage alone is sufficient to achieve watermark capacity in some situations. QIM performance of some relatively simple quantizers with a space filling gain and with a shape gain are compared and found to support the idea that space-filling gain is significant for QIM applications of quantizers.

The first section provides some notation and a brief review of QIM. Then the next section investigates the construction of vector quantizers for QIM that have a space-filling gain and no shape or memory gain. Final sections present QIM performance results, and the final section presents conclusions.

20

NOTATION AND REVIEW OF QIM

The basic watermarking task is to communicate a watermark message M by modifying a host signal X_1^N to produce the watermarked signal S_1^N . The watermarked signal then undergoes possible distortions over a watermark signal channel $p(Y_1^N | S_1^N)$ resulting in the received signal Y_1^N . In some cases it will be helpful to consider the noise introduced by the watermark signal channel as additive with variance σ^2 . Throughout this document, the host signal, watermarked signal, and received signal will be assumed to be taken from \mathbb{R}^N , an N-dimensional set of real numbers. An embedding rate R_m is used to embed the message in the host signal; accordingly, the message

$M \in \{0, 1, 2, \dots, 2^{NR_m} - 1\}$. It will be assumed that the distortion introduced by the watermark signal channel is independent of the message, and that the host signal is also independent of the message. The goal in the design of the watermarking system is to minimize the error probability in transmitting the message subject to a constraint D_m on the distortion introduced by the watermarked signal. The distortion measure used here will be the per letter squared error:

$$D(S_1^N, X_1^N) = \frac{1}{N} \sum_{i=1}^N (S_i - X_i)^2.$$

QIM Watermarking

Figure 1 shows how QIM can be used to implement the watermarking task. The watermark encoder associates a separate quantizer Q with each of the possible 2^{NR_m} messages. It chooses the quantizer Q_m corresponding to the input message M and uses this quantizer to quantize samples of the host signal, producing the watermarked signal S_1^N . These samples may be spatial, frequency, time or other transform domain features of the media signal. After this signal is passed through the channel, the watermark decoder's task is to estimate the message M based upon the received signal Y_1^N . This is done by quantizing Y_1^N using each of the 2^{NR_m} quantizers available to the encoder. The corresponding distortion is calculated for each quantizer, and the quantizer resulting in the lowest distortion is assumed to be the one most likely to have been used by the encoder. Then the decoder chooses the message \hat{M} , which corresponds to this quantizer, as its estimate of the original message M . The watermark encoder can be defined by

$$F(m, x) = Q_m(x),$$

and the decoder by

$$G(y) = \arg \min_{0 \leq m \leq 2^{NR_m}-1} D(y, Q_m(y))$$

Figure 2 illustrates dithered uniform scalar QIM. The message consists of a single bit that causes the encoder to choose between quantizer Q_0 , shown at top, and

quantizer Q_1 , shown below. The decoder effectively uses the quantizer shown at the bottom of the figure to decide between the two possible message values.

5

Distortion

Chen and Wornell propose a postprocessing technique for QIM that they call *distortion compensated* QIM. See B. Chen and G.W. Wornell, "Preprocessed and Postprocessed Quantization Index Modulation Methods for Digital Watermarking," in Security and Watermarking of Multimedia Contents II, P.W. Wong and E.J. Delp, eds.,

10 Proc. SPIE 3971, pp. 48-59, 2000. They associate a scale factor $\frac{\Delta}{\alpha}$ with each quantizer and use the parameter α to add part of the quantization error back to the watermarked source signal:

$$F_{DC}(m, x) = Q_m(x, \Delta/\alpha) + (1-\alpha)[x - Q_m(x, \Delta/\alpha)]$$

15 The parameter α is restricted to $[0, 1]$, ensuring that $D(y, Q_m(y, \Delta/\alpha))$ is independent of α . Chen and Wornell showed that the watermark performance can be optimized by choosing α so that

$$\alpha = \frac{DNR}{DNR+1},$$

Where $DNR = D_m/\sigma^2$.

20

DEVELOPMENT OF SPACE-FILLING QUANTIZERS FOR QIM

Two Dimensions

Figure 3 shows how dithered uniform scalar QIM can be extended to QIM using optimal uniform two dimensional quantizers. In the one dimensional case, the Figure shows a region of \mathbb{R}^1 (one-dimensional real number set) which is partitioned into two decision regions, one for each possible message value. Each decision region contains a code value at its centroid and is a polytope with minimal second moment about its centroid (in this case, trivially, an interval). By tiling \mathbb{R}^1 with the two decision regions and codevalues, it is possible to construct the quantizers shown in Figure 2.

The right side of Figure 3 shows a similar diagram in two dimensions. Here the optimal space-filling polytope is the hexagon; correspondingly, a region of \mathbb{R}^2 (two dimensional real number set) is partitioned into three hexagons, each with a code value at its centroid. Tiling \mathbb{R}^2 with this region results in a hexagonal lattice; each lattice point is assigned to one of three quantizers, and each of the three quantizers themselves have code values in a hexagonal sublattice of the original lattice. If code value C_{o_0} is chosen to lie on the origin, then the lattice of all code values consists of:

$$\{p : p = Um, m \in \mathbb{Z}^2\}, \text{ where } U = \begin{bmatrix} 0 & \sqrt{3}\Delta \\ 2\Delta & \Delta \end{bmatrix};$$

where \mathbb{Z} is two-dimensional set of integers.

The lattices of the three quantizer's codevalues are given by

$$Q_i = \{p : p = Vm + o_i, m \in \mathbb{Z}^2\}, \text{ where } V = \begin{bmatrix} -\sqrt{3}\Delta & \sqrt{3}\Delta \\ 3\Delta & 3\Delta \end{bmatrix}, o_0 = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, o_1 = \begin{bmatrix} 0 \\ 2\Delta \end{bmatrix}, o_2 = \begin{bmatrix} \sqrt{3}\Delta \\ \Delta \end{bmatrix}.$$

Higher Dimensions

To get an idea of the performance which can be achieved from large dimensional quantizers having only a space-filling gain, the previous two dimensional case can be

extended to an arbitrarily large number of dimensions n . To simplify the analysis, it will be convenient to apply some results of high resolution quantization theory, which assumes that quantizers have sufficient output levels so that the source probability distribution is essentially uniform within each quantization cell. The results of the
 5 analysis will be valid to the extent that the watermark distortion constraint D_m is small in comparison to the variance of the watermark source.

Here it will be assumed that the watermark signal channel is an Additive White Gaussian Noise (AWGN) channel with variance σ^2 . For large n the channel noise will, with high probability, lie within an n -dimensional sphere of radius $\sqrt{n\sigma^2}$. Thus if the
 10 decision regions of the watermark decoder are spheres of that size centered about the corresponding quantizer codevalues, the probability of watermark decoding error will approach zero with increasing n . Maximizing the watermark message rate for a given watermark distortion constraint is equivalent to solving an n -dimensional sphere packing problem similar to a problem that can be used to show the capacity of the Gaussian
 15 channel.

This is illustrated in Figure 4, in which one decoding sphere for each possible message value is packed within a larger sphere. \mathbb{R}^n (an n -dimensional space of real numbers) can be tiled with this larger sphere to produce the set of all decoding regions used by the decoder. Each quantizer of the watermark encoder has quantization cells
 20 which are spheres of the same size as the larger sphere in Figure 4. These quantization cells are packed tightly to fill \mathbb{R}^n . Each of the quantizers is identical except for an offset in \mathbb{R}^n corresponding to the location of the corresponding decoding sphere in Figure 4. The capacity of this idealized QIM is given by the ratio of the volumes of the smaller and larger n -dimensional spheres. The larger sphere must have a normalized second moment
 25 about its centroid equal to nD_m , the distortion constraint. This is given by

$$\frac{n}{n+2}r^2 = nD_m. \text{ For large } n, \text{ the radius of the larger sphere approaches } \sqrt{nD_m}. \text{ Since}$$

the volume of an n -dimensional sphere is $k_n r^n$, the maximum number of decoding spheres that can fit inside the larger sphere is:

$$\frac{k_n(nD_m)^{n/2}}{k_n(n\sigma^2)^{n/2}} = 2^{\frac{n}{2}\log\frac{D_m}{\sigma^2}}.$$

This results in an operational watermark capacity for this space-filling QIM of

$C = \frac{1}{2} \log \frac{D_m}{\sigma^2}$ bits per sample. If the source X_1^N is iid Gaussian, the watermark capacity is given by

5 $C = \frac{1}{2} \log \left(\frac{D_m}{\sigma^2} + 1 \right)$ bits per sample.

It is interesting to note that if the outer sphere with radius $\sqrt{n(D_m + \sigma^2)}$ shown in Figure 4 is packed with decoding spheres of radius $\sqrt{n\sigma^2}$, exactly enough decoding spheres fit to reach capacity; a QIM scheme using this larger sphere, however, would not satisfy the distortion constraint D_m .

10

Higher Dimension Distortion Compensated QIM

The capacity of a higher dimensional distortion compensated QIM can be analyzed in a similar manner to the previous case. The larger sphere available for packing of decoding spheres has a radius of

$$r_{l_{\arg e}} = \frac{\sqrt{nD_m}}{\alpha} = \frac{(D_m + \sigma^2)}{D_m} \sqrt{nD_m}.$$

The distortion compensated watermarked signal will lie within a sphere of radius

$$r_{DC} = (1 - \alpha)r_{l_{\arg e}} = \frac{\sigma^2 \sqrt{nD_m}}{D_m}$$

centered about each quantizer reconstruction value. Each decoding sphere contains the sum of the distortion compensated watermark signal and the channel noise. Since these two are independent, each decoding sphere has a radius of

$$r_{decode} = \sqrt{\left(\frac{\sigma^4}{D_m^2}\right)nD_m + n\sigma^2} = \sqrt{n\sigma^2 \left(\frac{\sigma^2 + D_m}{D_m}\right)}.$$

Taking the ratio of the volumes of the two spheres gives a capacity of $C = \frac{1}{2} \log \frac{D_m}{\sigma^2}$ bits per sample, exactly reaching the theoretical capacity.

Figure 5 compares the idealized high dimension space-filling QIM performance with and without distortion compensation.

5

PERFORMANCE EXAMPLES

To investigate how vector quantizers with different types of quantization gain perform in a QIM application, QIM using three types of quantizers were implemented. The first quantizer was the dithered scalar quantization proposed by Chen and Wornell.

10 The second was the hexagonal lattice quantizer described above. The final type of quantizer simulated for QIM was a structured vector quantizer (SVQ). In all cases the QIM used a blocklength of 32 samples in which either one of two messages (for dithered scalar QIM and SVQ QIM) or one of three messages (for 2-D hexagonal lattice QIM) was embedded. In the case of 2-D QIM, the same message was embedded in each of 16 pairs of samples in the 32 sample block.

15 The message can be used to represent an auxiliary data symbol (e.g., binary, M-ary, etc.). Groups of redundant or different symbols may be formed into a digital watermark message payload and embedded in the source media signal. To embed such a payload, the encoder maps the messages to corresponding samples, or blocks of samples, 20 in the source signal. These samples may be spatial, time, frequency, or some other transform domain sample of the source signal.

25 Examples of multi-dimensional quantizers that can be used for QIM watermarking include trellis coded quantizers and lattice vector quantizers, such as a hexagonal lattice quantizer. These multi-dimensional quantizers enable exploitation of the shape filling gain described above.

Structured Vector Quantizers and QIM

SVQ exhibits no space filling gain, but does take advantage of the available shape gain. It is able to achieve the performance of an entropy coded quantizer, but at a fixed

bit rate. The first step in designing an n -dimensional SVQ is to first design an N level entropy coded scalar quantizer and to calculate the codeword lengths corresponding to each of the quantizer's levels. Then an n -dimensional product codebook, containing N^n codevectors of length n is formed. Each vector in the product codebook is assigned a
 5 length corresponding to the sum of the lengths of the corresponding scalar quantizer codeword outputs. Then a threshold $T \approx nH(X)$, where $H(X)$ is the entropy of the source is chosen. Finally each codevector with a length greater than T is then removed from the product codebook. Source vectors of length n are then quantized to the closest codevector in the remaining codebook. The number of bits required to represent each
 10 codevector is approximately T , which is less than $n\log_2 N$. As block length increases, the performance of the SVQ approaches the distortion and rate performance of the original entropy coded scalar quantizer.

For QIM, two SVQs were designed, both with block lengths of 32 samples. The scalar quantizers used as the basis of these two SVQs were dithered uniform scalar
 15 quantizers. Logarithms of scalar quantizer probabilities were used instead of codeword lengths.

Results

Since the message size is different for the 2-D hexagonal lattice QIM than for the
 20 other two types of QIM, they cannot be directly compared by using the resulting probabilities of decoding error. Instead, the delivered channel capacities are derived from the decoding error probabilities and compared. The channel capacities are given by:

$$C_{DITHER} = 1.0 + P_{err} \log(P_{err}) + (1 - P_{err}) \log(1 - P_{err}),$$

$$C_{SVQ} = 1.0 + P_{err} \log(P_{err}) + (1 - P_{err}) \log(1 - P_{err}),$$

$$C_{HEX} = \log(3) + P_{err} \log(P_{err}/2) + (1 - P_{err}) \log(1 - P_{err}).$$

Figure 6 shows the results for a unit variance Gaussian source with AWGN of 22dB. The hexagonal lattice based QIM performed the best, followed by the dithered scalar QIM, and the SVQ QIM had the poorest performance.

The loss in performance between the dithered scalar quantizer and SVQ is most likely due to a slight decline in the distortion characteristics of the SVQ, when compared with the original scalar quantizer it was designed from. In a compression application, this decline is more than compensated for by the lower rate requirement of the SVQ
5 compared with the scalar quantizer. Since encoding rate is not important for QIM, this advantage does not accrue to the SVQ QIM.

The hexagonal lattice based QIM performed significantly better than the dithered scalar quantizer; in most cases the capacity of the hexagonal lattice QIM was approximately twice the capacity of the dithered scalar QIM.

10

CONCLUDING REMARKS

The results of the previous section have shown that it is possible to improve QIM performance by using a quantizer which provides a space-filling gain; the results have also shown that quantizers which provide a shape gain do not necessarily improve QIM performance. The capacity argument presented above implies that, asymptotically in
15 block length and quantizer resolution, space filling gain alone may be enough to achieve watermarking capacity in distortion compensated QIM.

Having described and illustrated the principles of the technology with reference to specific implementations in the attached paper and documents incorporated by reference,
20 it will be recognized that the technology can be implemented in many other, different, forms. To provide a comprehensive disclosure without unduly lengthening the specification, applicants incorporate by reference the patents and patent applications referenced above.

The methods, processes, and systems described above may be implemented in hardware, software or a combination of hardware and software. For example, the auxiliary data encoding processes may be implemented in a programmable computer or a special purpose digital circuit. Similarly, auxiliary data decoding may be implemented in software, firmware, hardware, or combinations of software, firmware and hardware. The methods and processes described above may be implemented in programs executed from
25

a system's memory (a computer readable medium, such as an electronic, optical or magnetic storage device).

The particular combinations of elements and features in the above-detailed embodiments are exemplary only; the interchanging and substitution of these teachings with other teachings in this and the incorporated-by-reference patents/applications are also contemplated.

RECORDED MAIL